FORMULATION AND SOLUTION OF MATRIX GAMES
WITHOUT UTILITY FUNCTIONS

JOHN V. HALL

FORMULATION AND SOLUTION OF MATRIX

GAMES WITHOUT UTILITY FUNCTIONS

*****

John V. Hall

FORMULATION AND SOLUTION OF MATRIX

GAMES WITHOUT UTILITY FUNCTIONS

by

John V. Hall

Lieutenant, United States Navy

Submitted in partial fulfillment of
the requirements for the degree of

MASTER OF SCIENCE

United States Naval Postgraduate School
Monterey, California

1 9 6 2

FORMULATION AND SOLUTION OF MATRIX

GAMES WITHOUT UTILITY FUNCTIONS

by

John V. Hall

This work is accepted as fulfilling

the thesis requirements for the degree of

MASTER OF SCIENCE

from the

United States Naval Postgraduate School

ABSTRACT

A brief resume of the role of linear utility functions in Game Theory is given. The point is made that, in practical applications, a knowledge of these functions is usually not available, and hence much of the rationale of the game theoretic approach to competitive problems is lost. The random character of the "real" payoff of a matrix game is then discussed, and the probability distribution function of the payoff is derived. The dependence of this distribution function upon the mixed strategies of the players is shown. Criteria are developed to provide definition of "optimal mixed strategy" in terms of the effect on the distribution function. The mathematical formulation of the solution is given for each criterion discussed.

The reader will require knowledge of the elements of Probability Theory, Game Theory, Utility Theory, and Linear Programming.

TABLE OF CONTENTS

1. Introduction.

The purpose of this paper is to investigate a difficulty that inevitably exists whenever an application of the mathematical Theory of Games is made in a non-idealized situation. We shall restrict our attention to the class of two-person matrix games, although the basic question arises in connection with n-person and continuous games as well. Briefly stated, the problem is as follows:

In the logic of the Theory of Games, a vital part is played by the so-called "linear" utility function, a real-valued function defined over the space of outcomes associated with a game. In actual practice, it is very doubtful that this function will be available to, let us say, a military operations analyst who wishes to formulate a practical problem as a game. On the other hand, when the situation involves a conflict of interest between intelligent antagonists, with an element of chance affecting the outcome of any course of action, it is difficult to find a concept of "rational procedure" or "optimal decision" which does not involve the game theoretic approach. Therefore the question becomes: in the absence of an important part of its machinery, what adaptations of Game Theory can provide us with reasonable solutions to problems of practical importance?

In standard texts on the Theory of Games, it is usual to present a discussion of utility theory at an early stage. The axioms leading to construction of the linear utility function (hereafter referred to as the l.u.f.) are developed, and once this has been done, it is customary to assume that the appropriate l.u.f.'s are known, and proceed with the mathematical development of the theory.

We wish to change this procedure slightly, in order to do two things:

(1) Clarify the theoretical role of the l.u.f. in the logic of Game Theory.

(2) Emphasize the vital importance to the military operations analyst of a knowledge of appropriate l.u.f.'s, when formulating "real" problems in game-theoretic terms.

2. Typical formulation of a competitive situation as a matrix game. Distribution of the payoff.

Suppose that a situation involving conflict of interest between two parties ("player I" and "player II") has been analysed so as to yield the following quantities:

(1) The space of all pure strategies available to player I.

(2) The space of all pure strategies available to player II.

(3) A matrix, the elements of which describe what happens when player I follows his i^{th} pure strategy, and player II follows his j^{th} pure strategy ($i = 1, 2, \ldots m$; $j = 1, 2, \ldots n$).

We shall assume that the outcome, or payoff to player I, resulting from the choice (i,j) of pure strategies, can be measured in terms of a single "real commodity", e.g., money, or casualties, etc.[1]

The first important point we make here is that nothing in Game Theory requires that a fixed real-commodity payoff to player I be associated with the pair (i,j). Rather, the requirement is this: the payoff to player I resulting from the choice (i,j) is a random variable X_{ij}, the distribution function of which is known to both players. Denote this distribution function by $F_{ij}(x) = P[X_{ij} \leq x]$. This leads us to the

___

1. This restriction simplifies the notation without affecting the logic of the argument.

2

following Definition: $F(x) = \left\lVert (F_{ij}(x)) \right\rVert$ is a matrix of known distribution functions. This matrix describes in probabilistic terms the payoff to player I associated with each possible choice (i,j) of pure strategies by the players.

Now assume that the players wish to employ mixed strategies.

$Q' = (q_1, q_2, \ldots q_m)$ is a mixed strategy of player I.

$P' = (p_1, p_2, \ldots p_n)$ is a mixed strategy of player II.

When the mixed strategies Q and P are used, we can determine the probability distribution function of the payoff to player I as follows:

Define Y to denote the real commodity payoff to player I when the mixed strategies Q and P are used in conjunction with the matrix $F(x)$. Then

$$G(y) \equiv P[Y \leq y] = \sum_{i,j} P[Y = X_{ij}] \; P[X_{ij} \leq y] \; , \; i.e.,$$

$$G(y) = \sum_{i,j} q_i p_j F_{ij}(y) \; . \tag{eq. 1}$$

This can be written in matrix notation as:

$$G(y) = Q'F(y)P \; . \tag{eq. 2}$$

Looking at equations 1 and 2, we can see that the distribution of the "real" payoff to player I involves two types of parameters: (1) those associated with the distributions $F_{ij}(x)$, which are assumed given and unchangeable, and (2) the elements q_i, p_j of the matrices Q and P. These are under the control of the players.

We can now formulate the objective of player I in choosing his mixed strategy Q. Player I wishes to choose Q so as to optimize, in some sense, the form of the distribution function $G(y)$. This interpretation of the role of a mixed strategy is a central feature of this paper.

At this point, the theory of the l.u.f. enters the picture. In the next section we see how the l.u.f. provides a criterion whereby each player may make a preferential ordering of his possible mixed strategies, and select an optimal one.

3.   The theoretical and practical roles of the l.u.f.

Appendix A presents for quick reference the axiomatic development of the l.u.f., largely following the excellent treatment in Luce and Raiffa [2] . Appendix B gives a very slight extension of the idea, primarily so that we can talk about a continuous l.u.f. u(x) defined over an "infinitely divisible" real commodity[2].

In this section, we do no more than write down some conclusions and interpretations concerning the l.u.f. Appendices A and B are, hopefully, arranged so that the reader who wishes to check the reasoning behind the statements made here may do so easily.

The pertinent conclusions concerning the l.u.f. are the following:

(1) The axiomatic design of a player's l.u.f. is such that it calls on information not available in the matrices $F(x)$, Q, and P, in order to measure the true value that a real commodity has for the player, vis-a-vis the gambles involved in a risky situation (i.e., the game).

(2) Suppose a player enters a gamble of the following form: he will win (or lose) a random amount X of some commodity, where X has the known distribution function H(x). This is called a "simple" gamble. If the player has a l.u.f. u(x) defined over the commodity, he can compute the true value v which the simple gamble has for him by forming the

2. The notation tends to be simpler when discussing this case.

4

expectation

$$v = E[u(X)] = \int_{-\infty}^{\infty} u(x)dH(x) \quad .$$

(3) Suppose the player is involved in a "compound" gamble; that is, with probability a_k he enters a simple gamble $H_k(x)$, k = 1,2,....n. The true value v to the player of this compound gamble can be computed in terms of his l.u.f. as[1]:

$$v = E[u(X)] = \int_{-\infty}^{\infty} u(x)d\left[\sum_k a_k H_k(x)\right] = \sum_k a_k \int_{-\infty}^{\infty} u(x)dH_k(x) \quad .$$

(4) Even though the operation involved in (2) and (3) above is that of taking an expected value of u(X), the interpretation of v as a "true value" does _not_ depend upon there being many repetitions of the gamble. The l.u.f. is so designed that it measures the value of a one-time gamble.

(5) Once a player has determined what his l.u.f. u(x) is for a certain commodity x, he can use u(x) to compare any two gambles involving x. If $H_1(x)$ and $H_2(x)$ are the distribution functions associated with any two such gambles, simple or compound, the player computes

$$v_1 = \int_{-\infty}^{\infty} u(x)dH_1(x) \qquad \text{and} \qquad v_2 = \int_{-\infty}^{\infty} u(x)dH_2(x) \quad .$$

The gamble involving $H_1(x)$ is preferred to the gamble involving $H_2(x)$ if and only if $v_1 > v_2$. By using u(x) in this way, the player obtains a complete and transitive ordering of his preference for all possible gambles over the commodity x.

These are the five conclusions which are of most interest to us. Next, we make some interpretations, based on these conclusions.

1. This property of u(x) contributes the unfortunate word "linear" to the name of the l.u.f.

5

(1)   When the players in a two-person matrix game make a choice of pure strategies, a "simple" gamble results. From player I's point of view, this gamble involves the random variable X_{ij} , with distribution function $F_{ij}(x)$.

(2)   When the players employ mixed strategies Q and P, player I faces a "compound" gamble involving the random variable Y with distribution function $G(y)$, as given by equation 1.

(3)   Recall our earlier observation that by his choice of a mixed strategy Q, player I exerts himself to determine a form of $G(y)$ which is in some sense optimal for him. We now interpret the theoretical role of the l.u.f. in the logic of Game Theory to be: the l.u.f. is a criterion-providing device that enables a player to compute which form of $G(y)$, and hence which mixed strategy, is best for him.

(4)   Suppose that a matrix game involves a real commodity x for which the players know their l.u.f.'s to be $u_1(x)$ and $u_2(x)$ respectively. Suppose further that $u_2(x) = -u_1(x)$. This is of course the so-called "zero-sum" assumption. If the players use mixed strategies Q and P, we write the distribution of the real payoff Y as $G(y) = G(y;Q,P)$ to emphasize the dependence of this function upon the players' choices. It now follows from the discussion under conclusion (5) above, that for extreme conservatism the optimal choices $Q_0' = (q_1^0,\ldots.q_m^0)$ and $P_0' = (p_1^0,\ldots.p_n^0)$ of mixed strategies are defined by the equation

$$\int_{-\infty}^{\infty} u(y)dG(y;Q_0,P_0) = \max_{Q} \min_{P} \int_{-\infty}^{\infty} u(y)dG(y;Q,P)$$

(eq.3)

$$= \min_{P} \max_{Q} \int_{-\infty}^{\infty} u(y)dG(y;Q,P) \quad .$$

Written in this form, the definition of optimal Q_0, P_0 looks unfamiliar

6

and rather formidable. However, referring to equation 1, we can re-write
equation 3 as:

$$\sum_{i,j} q_i^o p_j^o \int_{-\infty}^{\infty} u(y) dF_{ij}(y) \; = \; \max_{Q} \min_{P} \sum_{i,j} q_i p_j \int_{-\infty}^{\infty} u(y) dF_{ij}(y)$$

(eq. 4)

The "linear" property of $u(x)$ discussed under conclusion (3) above makes
it proper for us to define

$$a_{ij} \equiv \int_{-\infty}^{\infty} u(y) dF_{ij}(y),$$

and to re-write equation 4 as

$$\sum_{i,j} q_i^o p_j^o a_{ij} \; = \; \max_{Q} \min_{P} \sum_{i,j} q_i p_j a_{ij}.$$

(eq. 5)

By introducing the utility matrix $A \equiv \|(a_{ij})\|$, we can reduce equation 5
to the familiar form:

$$Q_o' A P_o \; = \; \max_{Q} \min_{P} Q'AP \; = \; \min_{P} \max_{Q} Q'AP \; .$$

(eq. 6)

Thus we have seen, step by step, how knowledge of a player's l.u.f.
for a commodity x lies at the center of the famous minimax formulation of
the zero-sum game solution, as given in equation 6. The essential obser-
vation to make is: if the l.u.f.'s are not known, the rationale leading
to equation 6 as a definition of optimality breaks down, for there can be
no guarantee that the players are maximizing (minimizing) the "true value"
of the compound gamble represented by the game. In particular, it is <u>not</u>
in general correct to replace a_{ij} by the quantity $m_{ij} \; = \; \int_{-\infty}^{\infty} x dF_{ij}(x)$, and
to maximize (minimize) the expected <u>real</u> value of the game. To do so
ignores the entire scheme for evaluating gambles that is provided by the

construction of the l.u.f.

This completes our review of the role of the l.u.f. and its practical importance to the operations analyst who wishes to use the "customary" approach to the solution of a matrix game, as represented in the zero-sum case by equation 6.

In the next section, we discuss procedures for handling a competitive situation when a knowledge of l.u.f.'s is not available. It seems appropriate, however, to first answer the question, "Why not devote the effort instead to the problem of determining the l.u.f. experimentally, rather than adopting the pessimistic attitude that appropriate l.u.f.'s are not likely to be known?"

The answer, of course, is that a great deal of work has been done in an effort to measure representative l.u.f.'s, and the results have been largely discouraging, not to say unbelievable. See, for instance, [3] and [4]. Since experiments performed under controlled and simplified laboratory conditions have failed to yield satisfactory results, it does not seem reasonable to assume that the military operations researcher, working with far more complex problems of value, will be able to isolate functions which truly describe how responsible military commanders evaluate difficult situations in terms of "linear utilities". As a practical matter, it seems much more sensible to look for procedures which make some sense in the absence of knowledge about linear utilities.

One more point should be made before we go on. If we are going to abandon all knowledge of utility functions and attempt to "solve" competitive situations (games) using only the distribution function $G(y)$ as a point of departure, we must clearly understand what we can and cannot hope to accomplish.

8

Consider the situation in which a matrix game has been formulated, and a knowledge of appropriate l.u.f.'s assumed. In this case, if the payoff is zero-sum, the assertion of Game Theory is that there is a logical and highly plausible definition of "optimum solution" available, and a scheme for finding it[1]. In other words, we are guaranteed the optimum result.

Contrast this with the situation that exists when no knowledge of the l.u.f. is assumed. This lack of knowledge means that we do not know how to value the commodity for which we are gambling, vis-a-vis the risks involved. It is clearly unrealistic to hope that, in this condition of ignorance, we shall be able to characterize any decision we make as the optimal one. We can hope to formulate criteria which may be "plausible" and "reasonable" when applied to specific situations, and we can hope that these criteria, through their mathematical formulation, will lead to decisions which are "approximately optimal". This is of course a very loose statement, since we do not have, and will not formulate, a means of measuring "distances" between a theoretical optimum and any other point in the space of possible solutions. We are simply proceeding in accordance with the belief that it is better to have some well-understood criterion for the goodness of a decision, and some systematic scheme for computing what that decision should be, keeping in mind all the potential shortcomings, than to have none at all and simply proceed hit-or-miss.

4. Criteria for the solution of matrix games in the absence of linear utility functions.

Consider first the case of a game that is going to be played one time only. For this case, we shall formulate solutions that are approximately

1. We have not discussed such schemes so far, but the "simplex method" may be regarded as a general algorithm for solving this class of games.

optimal under the following criteria:

Criterion A: That mixed strategy is best for player I which maximizes the probability that the real-commodity payoff will lie in any one of a finite number of given, disjoint intervals, assuming that player II acts so as to minimize this probability.

Criterion B: That mixed strategy is best for player I which

i) maximizes the probability that the real-commodity payoff will exceed a given amount b, subject to the condition that

ii) the probability that the payoff is less than a given amount c, shall be less than a given figure a, no matter what strategy player II uses. Assume that player II wishes to minimize the probability that player I receives more than b.

These are certainly not the only criteria one could consider. However, they have features which are of practical value. Three obvious special cases of Criterion A call for

(1) maximizing the probability that the real payoff exceeds a given value (perhaps some "critical" value),

(2) minimizing the same probability,

(3) maximizing the probability that the payoff lies in some specified interval.

Certainly one can think of practical situations in which one of these might be a very reasonable course of action to follow. What is more important, it is not hard to imagine that a military commander might have just enough information to decide on one or two "critical" values, when a utility

10

measure over all possible outcomes would be impossible to find. As an example, the commander who decides, "I must at all costs destroy at least 80% of the enemy fleet in this engagement", has adopted special case (1) above as his criterion for an optimal course of action. A more homely example would be an expert poker player who decides, "I want to win at least twenty dollars in this game, but I'd better not win more than fifty, or these people might not let me fleece them in the future." The poker expert is looking toward special case (3).

Obviously, Criterion A is worth considering. The general formulation is no more difficult than the special cases, as will be seen[1].

It is a little harder to read sense into Criterion B at first glance, but an example will make it clear that it too has a plausible, common-sense basis. Suppose a businessman has the following thoughts: "I would like very much to make a profit of $100,000 or more this year. ($100,000 = given amount b). On the other hand, I'll have to take some risks to do that. I have some debts outstanding, my son enters Harvard this September, and I want to vacation in Hawaii this year. If I show a profit of less than $10,000, things are going to be tough. ($10,000 = given amount c). Well, I'll try for the big profit, but only if I can work it so that my chances of making less than $10,000 are small; say about 10%." (10% = given figure a).

One can easily construct examples with a military flavor, involving the value of an objective, and the risks and costs associated with achieving it. As with criterion A, the point is this: It is far more reasonable to expect our businessman in the example to somehow get hold of the figures $100,000, $10,000, and 10%, than it is to expect him to construct his

1. It is also convenient to solve this problem for any finite number of intervals, since solutions for all the more practical special cases follow immediately.

complete l.u.f. for money. As for asking the Joint Chiefs of Staff what
their combined l.u.f. for casualties in Southeast Asia might be ---one
shudders at the thought.

In connection with both Criteria, we have made the assumption, essent-
ially, that player II's aims will be directly opposed to those of player I.
Obviously, in many real situations this will not be true. When we deal
with this difficulty, we get into the whole problem of non-zero-sum games,
definitions of stable solutions, etc.

Now, in the absence of l.u.f.'s, or any measure of "true value" assoc-
iated with real outcomes of a game, we lack the means for deciding whether
or not a game is zero-sum. Hence, we should really consider the difficult
non-zero-sum case. To get around this situation, we adopt the following
point of view, which should be regarded as part of Criteria A and B, and
therefore as an "input" to any solution stemming from them:

The most conservative policy that player I can follow is to assume
that player II is an "implacable foe". In practical situations where this
is largely true, all will be well, and incorporation of this assumption
will be largely justified. When this assumption is not true, player I will
get at least slightly better results than if player II were "implacably
opposed", but will not be able to take maximum advantage of possible
"community of interest" arrangements, etc. In this latter case, the pro-
cedures which we develop here probably will not be very good, and player I
should exercise his ingenuity to analyze the special features of the part-
icular situation, in order to find a better solution. In other words, a
practical application of these ideas requires a realistic look at the con-
text in which they are to be applied, and no good solution can be expected
from their use in inappropriate circumstances.

12

5. Mathematical formulation of the solution under criterion A.

We are still dealing with the case where the real payoff is one-dimensional, that is, expressible in terms of a single real commodity x.

Following the previous notation, let Q and P be mixed strategies of players I and II respectively, F(x) the matrix defined on page 3, and Y the real payoff to player I resulting from choices of Q and P. Denote the disjoint intervals of Criterion A by $I_1 = [a_1, b_1]$, $I_2 = [a_2, b_2]$, ..., $I_n = [a_n, b_n]$. Then

$$P\left[Y \in I_1 \text{ or } Y \in I_2 \text{ or} \ldots Y \in I_n\right] = \sum_{k=1}^{n} \int_{a_k}^{b_k} dG(y)$$

$$= \sum_{k=1}^{n} \int_{a_k}^{b_k} d\left[\sum_{i,j} q_i p_j F_{ij}(y)\right] \quad \text{by equation 1.}$$

$$= \sum_{k=1}^{n} \sum_{i,j} q_i p_j \int_{a_k}^{b_k} dF_{ij}(y) = \sum_{k=1}^{n} \sum_{i,j} q_i p_j \left[F_{ij}(b_k) - F_{ij}(a_k)\right].$$

Using matrix notation (see equation 2), re-write this as

$$Q'\left[F(b_1) - F(a_1)\right]P + Q'\left[F(b_2) - F(a_2)\right]P + \ldots + Q'\left[F(b_n) - F(a_n)\right]P$$

$$= Q'\left[F(b_1) + \ldots + F(b_n) - F(a_1) - \ldots - F(a_n)\right]P.$$

In the last expression, denote the matrix $\left[F(b_1) + \ldots - F(a_n)\right]$ by H. This is a matrix of real numbers, completely determined by F(x) and the intervals I_i, which are given quantities. Therefore, Q'HP is an ordinary bilinear form, and under our assumption about the behaviour of player II, conservative optimal mixed strategies Q_0 and P_0 are defined by

$$Q_0'HP_0 = \max_{Q} \min_{P} Q'HP. \quad \text{(eq. 7)}$$

13

The nice feature of this result is of course that all the machinery for finding Q_o and P_o that one uses in connection with zero-sum games, is appropriately used here. The difference lies in the interpretation of H as a matrix whose elements come from probability calculations, rather than a utility function.

In the special case of maximizing $P\left[Y > c\right]$ where c is some "critical" value, the result is

$$P\left[Y > c\right] = \int_c^{\infty} dG(y) = \sum_{i,j} \int_c^{\infty} dF_{ij}(y) = \sum_{i,j} q_i p_j \left[1 - F_{ij}(c)\right] \quad .$$

To put this in the more convenient matrix notation, define 1_{mn} to be a m-by-n matrix each element of which is unity, and write

$$P\left[Y > c\right] = Q'\left[1_{mn} - F(c)\right] P \quad .$$

Optimal strategies Q_o and P_o are defined by

$$Q_o'\left[1_{mn} - F(c)\right] P_o = \max_Q \min_P Q'\left[1_{mn} - F(c)\right] P \qquad \text{(eq. 8)}$$

An interesting sub-special case develops when the probability distributions $F_{ij}(x)$ of the matrix $F(x)$ are degenerate, that is:

$$F_{ij}(x) = 0 \quad \text{for} \quad x \leq h_{ij}$$
$$\qquad\qquad = 1 \quad \text{for} \quad x > h_{ij} \quad .$$

This corresponds to the case where the real-commodity payoff associated with the choice (i,j) of pure strategies is a fixed amount h_{ij}, rather than random. In this case the matrix $\left[1_{mn} - F(c)\right]$ of equation 8 is a matrix in which each element is either zero or one. Choice of a simplified criterion certainly leads in this case to a simplified problem to be solved!

14

6. Mathematical formulation of the solution under Criterion B.

The procedure for finding an optimal strategy under Criterion B is somewhat more complicated than it is under Criterion A. We use a linear programming approach [5], [6]. The problem turns out to be of the "maximization" variety, which entails

(1) maximizing a linear form $\sum_i w_i t_i$ in non-negative variables t_i, subject to

(2) a set of linear constraints (inequalities) on the variables t_i.

Once such a problem has been correctly formulated, the set[1] of values of the variables t_i that maximizes the linear form may be found (provided a set satisfying the constraints exists) by the "Simplex Method" [5], [6], which may be regarded as a general algorithm for solving linear programming problems. In any particular case, the amount of computation involved may be enormous, so that a high-speed computer may be needed to get actual numerical results. Conceptually however, the solution is always computable when it exists.

In Appendix 5 of their book "Games and Decisions", Luce and Raiffa give an excellent formulation of the ordinary two-person zero-sum game solution as a linear programming problem. The problem here turns out to be of a very similar form, so the reader might wish to refer to Luce and Raiffa before going on.

Using the same notation as before, the solution is developed in nine steps as follows:

(1) Let $F(b)$ have dimensions m x n, and consider the matrix equation

$$F(b)' \ Q \ = \ B \ .$$

1. Not necessarily unique.

The n x 1 matrix B is _defined_ by this equation, and consists of elements B_j $(j = 1,2,..,n)$ with

$$B_j = \sum_{i=1}^{m} q_i F_{ij}(b) .$$

It is easy to see that B_j is the probability that $Y \le b$, when player I uses mixed strategy Q and player II uses his j^{th} pure strategy. Note that $B_j \ge 0$ for all j.

    (2) Define $B_{jo} \equiv \max_{j} B_j$.

Player I wishes to maximize $P\left[Y > b\right]$, and therefore to minimize $P\left[Y \le b\right]$. Hence, under our assumption about the behavior of player II which was previously stated, the objective of player I becomes: choose Q so as to minimize B_{jo}.

    (3) Define the m x 1 matrix $U = (1/B_{jo})Q$, and define 1_n to be an n x 1 matrix, each element of which is unity. Then we have

$$(1/B_{jo})B \le 1_n ,$$

and consequently,

$$F(b)'U \le 1_n .$$

    (4) We can denote the matrix U by $U' = (u_1, u_2, .., u_m)$, where

$$u_i = (q_i/B_{jo}) , \quad i = 1,2,...,m .$$

Note that $u_i \ge 0$ for all i. The u_i are the non-negative variables of this linear programming problem.

16

(5) Next, note that

$$\sum_{i=1}^{m} u_i = \sum_{i=1}^{m} (q_i/B_{j_0}) = 1/B_{j_0} .$$

This can also be written as

$$1_m' U = 1/B_{j_0} ,$$

where $1_m'$ is an $m \times 1$ matrix of one's. Since player I wishes to minimize B_{j_0}, he therefore will maximize $1/B_{j_0}$, which is equivalent to maximizing the linear form $1_m' U$.

(6) Thus, $1_m' U$ is the linear form in non-negative variables to be maximized, and from step (3) above, the matrix equation $F(b)'U \leq 1_n$ represents a set of linear constraints.

(7) There is, however, another set of linear constraints to be considered. Recall that under Criterion B, we demand $P[Y \leq c] \leq a$, against each pure strategy of player II. The matrix equation which expresses this requirement is

$$F(c)'Q \leq a \ 1_n .$$

We need to include this equation among the constraints. To do this, recall that

$$1/B_{j_0} = 1_m' U .$$

Therefore,

$$F(c)'Q(1/B_{j_0}) \leq a \ 1_n (1_m'U) ,$$

or

$$F(c)' U \leq a \ 1_{mn}' U ,$$

where 1_{mn} is an $m \times n$ matrix, each element of which is unity. Now define 0_n to be an $n \times 1$ matrix of zeros, and rewrite the last equation as

$$\left[F(c)' - a\ 1'_{mn} \right] U \leq 0_n \quad.$$

(8) Thus far, then, we have the problem

maximize $1'_m\ U$,

subject to

$$F(b)'U \leq 1_n \quad \text{and} \quad \left[F(c)' - a\ 1'_{mn} \right] U \leq 0_n \quad.$$

The two equations expressing the constraints can be combined into one

matrix equation. Define two partitioned matrices A and C as follows:

$$A' \equiv \left\| \begin{array}{c} F(b)' \\ \text{------------} \\ F(c)' - a\ 1'_{mn} \end{array} \right\| \quad \text{and} \quad C \equiv \left\| \begin{array}{c} 1_n \\ \text{----} \\ 0_n \end{array} \right\| \quad.$$

(9) The linear programming problem may now be written in final form

as

maximize $1'_m\ U$, subject to $A'\ U \leq C$.

Of course, once U has been determined, player I's optimal strategy

is found by computing $1/B_{jo} = \sum_i u_i$, so that $Q_{optimal} = B_{jo}\ U$.

In conclusion, note that under Criterion B all the information neces-

sary to construct the matrix A is assumed given, so that only the comput-

ational problem remains. Furthermore, one can determine whether or not

a solution exists by computing $\min_Q \max_P Q'\ F(c)\ P$. If this quantity is

greater than a, it is obvious that there can be no solution; conversely,

if it is smaller than a, there must be a solution.

7. <u>Extension to the case of a real payoff which is not one-dimensional.</u>

In the previous sections we considered games whose outcomes were measur-

able in terms of a single real commodity, e.g., money. Let us now consider

a case which, if not completely general, surely possesses enough general-

ity to be of practical use.

Suppose that a game can result in any one of r different outcomes, or "states" as we shall call them. Denote these states by S_j , $j = 1, 2, \ldots, r$. S_j must be distinguishable from S_i , for all $i \neq j$.

Next, suppose that player I can make a complete and transitive ordering of the states S_j in terms of preference, where the word is used in its everyday sense. By this we mean

(1) For any two states S_i and S_j , $i \neq j$, player I either prefers S_i to S_j , or vice versa.

(2) If player I prefers S_i to S_j, and S_j to S_k , then he prefers S_i to S_k .

When player I is indifferent between S_i and S_j, i.e., has equal preference for them, we interpret this to mean that the relations "S_i preferred to S_j" and "S_j preferred to S_i" hold simultaneously. However, in order to simplify notation, we shall assume a strict ordering exists, with no cases of indifference. This assumption does not affect the validity of what follows; the case of indifference can be included at the expense of some extra symbols.

Now, let the states be re-numbered in ascending order of preference, that is, S_1 is the least preferred state, S_2 the next-to-least preferred, and so on until S_r is the most preferred. During this process, we retain the identity of each state, that is, we still know what real outcome each of the symbols S_j represents.

Next we define S to be the ordered set $S = \{S_1 , S_2 , \ldots, S_r\}$ of the re-numbered states. It is important to note that the elements of S are disjoint.

The purpose of all this is to lead to the following function. We define X as a function with domain S and range the set of integers

19

$\{1,2,\ldots,r\}$, such that

$$S_j \xrightarrow{X} j \qquad \text{for all } S_j \in S.$$

The function X is of course a random variable defined on the "sample space" S. Inasmuch as the probability distribution over S is known[1] for each choice (i,j) of pure strategies, the distribution of X is also known. Let us illustrate this fact. Suppose that corresponding to the choice (i,j) of pure strategies, the distribution of the outcome was

$$P_{ij} \left[\text{outcome is } S_k \right] = p_k \qquad \text{for all } S_k \in S$$

$$= 0 \qquad \text{elsewhere.}$$

The corresponding distribution of X is given by

$$P_{ij} \left[X = k \right] = p_k \qquad \text{for } k = 1, 2, \ldots, r$$

$$= 0 \qquad \text{elsewhere.}$$

Hence, given the matrix $F(S_k)$, $k = 1,2,\ldots,r$, of distribution functions (which corresponds to the matrix F(x) of previous sections), we can immediately construct the matrix $F^*(k)$, $k = 1,2,\ldots,r$, the elements of which are the distribution functions of X, corresponding to the pairs (i,j) of pure strategies.

Since we are dealing with a discrete random variable, we can also construct the matrix of probability mass functions which corresponds to $F^*(k)$. Denote this matrix of mass functions by J(k).

Now we can repeat essentially all the computations of previous sections. For instance, if Q and P are the mixed strategies of the players, then (in terms of player I's preferences)

1. By hypothesis; see p. 2.

20

P [outcome is preferred to S_j] = P [outcome is S_{j+1} or

S_{j+2} or...or S_r] = P [X = j+1] + P[X = j+2] + ... + P[X = r]

= Q' J(j+1) P + Q' J(j+2) P+ ...+ Q' J(r) P

= Q' [J(j+1) + J(j+2) + ... + J(r)] P .

Another example:

P [outcome is either S_i or S_j or...,or S_k] =

P[X = i] + P[X = j] +....+ P[X = k] =

Q' [J(i) + J(j) + ...+ J(k)] P .

We shall not pursue the details, but it should be clear, and the reader
can easily verify, that everything goes through as before, including the
solutions under Criteria A and B.

Note that we have said nothing about the nature of the states S_j.
They may be defined in any way at all, quantitatively or qualitatively.
We have required only that they be distinguishable, and that player I must
have a transitive preference ordering over them.

8. Some comments concerning the distribution function G(y).

In section 2 we derived the distribution function

$$G(y) = \sum_{i,j} q_i P_j F_{ij}(y) = Q'F(y)P .$$

It seems worthwhile to point out that G(y) represents a special case
of an interesting class of distribution functions called "mixture distrib-
utions". For an excellent discussion of this type of function, and some
of its history, see [7]. Generally speaking, a distribution of this type
arises when a sample is to be drawn from one of several populations, but
only a probability statement can be made as to exactly which one.

To illustrate one of the interesting properties of mixture distributions,

21

let us find the mean of $G(y)$.

$$E(Y) = \int_{-\infty}^{\infty} y dG(y) = \sum_{i,j} q_i P_j \int_{-\infty}^{\infty} y dF_{ij}(y) = \sum_{i,j} q_i P_j m_{ij},$$

where $m_{ij} = E(X_{ij})$. If $M = \left\| (m_{ij}) \right\|$, then

$$E(Y) = Q'MP.$$

It is easy to show that, if $M^n = \left\| (E(X_{ij}^n)) \right\|$, then

$$E(Y^n) = Q'M^nP.$$

Consequently, the variance of $G(y)$ is given by

$$Var(Y) = E(Y^2) - E^2(Y) = Q'M^2P - m^2.$$

One may wish to ask, "If the distribution of Y can be written in terms of the distributions F_{ij} as shown, what is the functional relationship between Y and the X_{ij}?". This can lead to some confusion unless it is realized that Y is __defined__ by the equation $G(y) = P[Y \le y] = Q'F(y)P$. No other relation between Y and the X_{ij} is needed.

9. __Conclusion__.

There are two major conclusions which, it is hoped, the reader will have drawn from this paper. They are

(1) A major __practical__ stumbling block in the application of Game Theory to real problems is the unavailability of linear utility functions. We have seen how necessary a knowledge of these functions is, when it comes to making actual computations in order to solve a matrix game in the "usual" sense.

(2) Despite this difficulty, there are reasonable procedures which can

22

be adopted, provided circumstances warrant, in order to define and compute "optimal" solutions.

This paper is intended only to illustrate some ideas, and not to make a complete analysis of any particular class of solutions. We hope the reader will agree that the general scheme discussed here has possibilities for use in many situations other than the ones we have looked at.

BIBLIOGRAPHY

1. J. von Neumann and O. Morgenstern, <u>Theory of Games and Economic Behavior</u>, second edition, Princeton University Press, 1947.

2. R. D. Luce and H. Raiffa, <u>Games and Decisions</u>, John Wiley and Sons, 1957.

3. D. Davidson, P. Suppes, and S. Siegel, <u>Decision Making, an Experimental Approach</u>, Stanford University Press, 1957

4. <u>Decision Processes</u>, ed. by R. M. Thrall et. al, John Wiley and Sons, 1954.

5. <u>Activity Analysis of Production and Allocation</u>, Cowles Commission for Research in Economics, Monograph #13, ed. by T. C. Koopmans, John Wiley and Sons, 1951

6. S. I. Gass, <u>Linear Programming</u>, McGraw-Hill, 1958

7. LT E. A. Parent, (SC) USN, <u>Mixtures of Distributions</u>, U. S. Naval Postgraduate School, 1962

APPENDIX A

AXIOMATIC DEVELOPMENT OF THE LINEAR UTILITY FUNCTION

Chapter 2 of "Games and Decisions" by Luce and Raiffa gives an excellent presentation of this subject, and the reader is referred to their book for a detailed discussion of the axioms which follow.

Definitions and notation. Assume all gambles or "lotteries" under discussion involve a finite set of prizes or alternatives, $A = \{A_1, A_2, \ldots, A_r\}$. A simple lottery is defined to be a chance mechanism which associates with each alternative A_i a known probability $p_i = P_L(A_i)$, $i = 1, 2, \ldots, r$, $\sum_i p_i = 1$. A simple lottery L is thus a probability distribution over A, and may also be regarded as an r-tuple, $L = \{p_1 A_1, p_2 A_2, \ldots, p_r A_r\}$.

A compound lottery is a two-stage chance mechanism, which, with probability q_i, yields a simple lottery $L^{(i)}$ as the result of its first stage. ($i = 1, 2, \ldots, s$; $\sum_i q_i = 1$). The simple lottery $L^{(i)}$ then constitutes the second stage. We symbolize a compound lottery by $\mathcal{L} = \{q_1 L^{(1)}, \ldots, q_s L^{(s)}\}$, an s-tuple. Note that $\mathcal{L}$ is again no more than a probability distribution over A, where $P_{\mathcal{L}}(A_i) = \sum_{j=1}^{s} q_j p_i^{(j)}$.

The purpose is to derive the utility function of an individual. When we speak of preferences, they are the preferences of that individual.

The symbol $\succeq$ expresses a preference ordering. That is, $A_i \succeq A_j$ means that the individual prefers A_i to A_j or is indifferent between them, and $L \succeq L'$ or $\mathcal{L} \succeq \mathcal{L}'$ expresses the same relationship between two lotteries. If $A_i \succeq A_j$ and $A_j \succeq A_i$ hold simultaneously, we write $A_i \sim A_j$. The symbol $\sim$ denotes indifference between the two alternatives.

Throughout what follows, it is assumed that the set A has been ordered so that A_1 is the most preferred alternative, and so on down,

25

until A_r is the least preferred. This can be done with no loss of generality.

Axiom 1. The preference ordering $\succsim$, a binary relation, yields a complete and transitive ordering of A.

"Complete" means, for any A_i and A_j in A, either $A_i \succsim A_j$ or $A_j \succsim A_i$.

Axiom 2. Let $\mathcal{L}$ be a compound lottery and L a simple lottery. If $P_L(A_i) = P_{\mathcal{L}}(A_i)$ for all A_i in A, then $\mathcal{L} \sim L$.

This says that it is only the character of a lottery _as a probability distribution_ which interests the individual.

Axiom 3. Each alternative A_i is indifferent to some simple lottery which mixes only A_1, the best alternative, and A_r, the worst alternative. That is, for every A_i in A, there exists some probability u_i, such that $A_i \sim [u_i A_1, \ (1 - u_i) A_r]$.

Axiom 4. If $A_i \sim [u_i A_1, \ (1 - u_i) A_r]$, then in any simple lottery $L = (p_1 A_1, \ \ldots, \ p_i A_i, \ \ldots, \ p_r A_r)$, $p_i A_i$ may be replaced by $p_i [u_i A_1, \ (1 - u_i) A_r]$.

Axiom 5. The relations $\succsim$ and $\sim$ are transitive among lotteries, just as they are among the elements of A.

This is obviously desirable, but it cannot be deduced from the preceding axioms, and so must be included separately.

Axiom 6. The relation $[pA_1, \ (1 - p)A_r] \succsim [p'A_1, \ (1 - p')A_r]$ holds if and only if $p \geq p'$.

These six axioms give us the following theorem, which is the point

26

of the entire development:

Theorem. If the relation $\succeq$ satisfies axioms 1 through 6, then there exist numbers u_i associated with the elements A_i of A, such that for any two lotteries L and L', the relation $L \succeq L'$ holds if and only if

$$\sum_{i=1}^{r} u_i P_L(A_i) \qquad \sum_{i=1}^{r} u_i P_L'(A_i) \quad .$$

The proof is obvious from the axioms. Any compound lottery can be reduced to a simple lottery involving $A_1,\ldots, A_r$, which can in turn be converted to another compound lottery involving just A_1 and A_r, which can be reduced to a simple lottery involving A_1 and A_r, which can by axiom 6 be compared with any lottery so treated. The "experimental" part of the program lies in finding the numbers u_i of Axiom 3. These first appear in the guise of probabilities, and later turn out to be the utilities which we sought. It is easy to show that any positive linear transformation of these numbers again yields a utility function which reflects the individual's preferences.

The reader who has followed through this derivation should easily be able to verify the assertions made in section 3 of this paper.

APPENDIX B

A NON-RIGOROUS EXTENSION OF THE LINEAR UTILITY CONCEPT.

In Appendix A, the assumption was made that the set A of alternatives associated with any lottery L was finite, and that the lottery comprised a probability mass function $P_L(A_i)$ defined on A. Suppose now we wish to enter a gamble G involving some numerically measurable and infinitely divisible commodity x. Associated with G there is a probability density function $f_G(x)$, defined over x.

We want to make all the logic and conclusions of Appendix A valid for this case also, so we take the following position. All "continuous" gambles discussed in this paper are such that

(1) If I is the interval over which x is defined, I may be partitioned into a finite number of disjoint subintervals dx_i, whose union is I. If X is the random variable of the gamble, $P[X \in dx_i] = f_G(x_i)dx_i$, where $x_i \in dx_i$.

(2) By proper choice of dx_i and x_i, we can approximate $f_G(x)$ as closely as we like by a probability mass function $f_G(x_i)dx_i$.

Once this has been done, all the logic of Appendix A will go through, and the numbers u_i can be found. We now <u>assume</u> that the set of numbers u_i thus produced can be represented, as $dx_i \rightarrow 0$ and the number of subintervals increases without limit, by a continuous curve u(x). When this is the case, it is intuitively clear that the continuous expected value calculation will have the same characteristics as the discrete calculation in Appendix A.